Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	43	(lin near derrick).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29
L2	20556	available near2 bandwidth	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29
L3	1250086	monitor$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29
L4	20556	available near2 bandwidth	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29
L5	1250086	monitor$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29
L6	1319	L4 same L5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29
L7	11387	register adj file	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29
L8	21856	memory adj (bus or channel)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29
L9	1319	L4 same L5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29

L10	21856	memory adj (bus or channel)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29
L11	74	L9 and L10	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29
L12	11387	register adj file	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29
L13	74	L9 and L10	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29
L14	10	L12 and L13	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29
L15	43	(lin near derrick).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29
L16	10	L12 and L13	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29
L17	2	L16 and L15	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/05/31 09:29

Terms used
available bandwidth; bus monitoring memory channel load capacity

Found **972** of **155,867**

Sort results by relevance

Display results expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Relevance scale

1 Cache Memories

Alan Jay Smith

September 1982 **ACM Computing Surveys (CSUR)**, Volume 14 Issue 3

Full text available: pdf(4.61 MB) Additional Information: full citation, references, citings, index terms

2 Bandwidth: System capability effects on algorithms for network bandwidth measurement

Guojun Jin, Brian L. Tierney

October 2003 **Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement**

Full text available: pdf(254.09 KB) Additional Information: full citation, abstract, references, citings, index terms

A large number of tools that attempt to estimate network capacity and available bandwidth use algorithms that are based on measuring packet inter-arrival time. However in recent years network bandwidth has become faster than system input/output (I/O) bandwidth. This means that it is getting harder and harder to estimate capacity and available bandwidth using these techniques. This paper examines the current bandwidth measurement and estimation algorithms, and presents an analysis of how these al ...

Keywords: algorithm, bandwidth, design, estimation, measure, network, performance, system capability

3 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**

Full text available: pdf(4.21 MB) Additional Information: full citation, abstract, references, index terms

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display

repeated occurrences of non-trivial commun ...

4 The evolution of the DECsystem 10

C. G. Bell, A. Kotok, T. N. Hastings, R. Hill

January 1978 **Communications of the ACM**, Volume 21 Issue 1

Full text available: pdf(1.92 MB) Additional Information: full citation, abstract, references, citings, index terms

The DECsystem 10, also known as the PDP-10, evolved from the PDP-6 (circa 1963) over five generations of implementations to presently include systems covering a price range of five to one. The origin and evolution of the hardware, operating system, and languages are described in terms of technological change, user requirements, and user developments. The PDP-10's contributions to computing technology include: accelerating the transition from batch oriented to time sharing computing systems; ...

Keywords: architecture, computer structures, operating system, timesharing

5 Local networks

William Stallings

March 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 1

Full text available: pdf(3.01 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The rapidly evolving field of local network technology has produced a steady stream of local network products in recent years. The IEEE 802 standards that are now taking shape, because of their complexity, do little to narrow the range of alternative technical approaches and at the same time encourage more vendors into the field. The purpose of this paper is to present a systematic, organized overview of the alternative architectures for and design approaches to local networks.

...

6 System-level power optimization: techniques and tools

Luca Benini, Giovanni de Micheli

April 2000 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 5 Issue 2

Full text available: pdf(385.22 KB) Additional Information: full citation, abstract, references, citings, index terms

This tutorial surveys design methods for energy-efficient system-level design. We consider electronic sytems consisting of a hardware platform and software layers. We consider the three major constituents of hardware that consume energy, namely computation, communication, and storage units, and we review methods of reducing their energy consumption. We also study models for analyzing the energy cost of software, and methods for energy-efficient software design and compilation. This survery ...

7 Separated high-bandwidth and low-latency communication in the cluster interconnect Clint

Hans Eberle, Nils Gura

November 2002 **Proceedings of the 2002 ACM/IEEE conference on Supercomputing**

Full text available: pdf(235.04 KB) Additional Information: full citation, abstract, references, index terms

An interconnect for a high-performance cluster has to be optimized in respect to both high throughput and low latency. To avoid the tradeoff between throughput and latency, the cluster interconnect Clint has a segregated architecture that provides two physically separate transmission channels: A *bulk channel* optimized for high-bandwidth traffic and a

quick channel optimized for low-latency traffic. Different scheduling strategies are applied. The bulk channel uses a scheduler that ...

8 RAID: high-performance, reliable secondary storage

Peter M. Chen, Edward K. Lee, Garth A. Gibson, Randy H. Katz, David A. Patterson
June 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 2

Full text available: pdf(3.60 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Disk arrays were proposed in the 1980s as a way to use parallelism between multiple disks to improve aggregate I/O performance. Today they appear in the product lines of most major computer manufacturers. This article gives a comprehensive overview of disk arrays and provides a framework in which to organize current and future work. First, the article introduces disk technology and reviews the driving forces that have popularized disk arrays: performance and reliability. It discusses the tw ...

Keywords: RAID, disk array, parallel I/O, redundancy, storage, striping

9 Bridging the digital divide: storage media + postal network = generic high-bandwidth communication

Nitin Garg, Sumeet Sobti, Junwen Lai, Fengzhou Zheng, Kai Li, Randolph Y. Wang, Arvind Krishnamurthy
May 2005 **ACM Transactions on Storage (TOS)**, Volume 1 Issue 2

Full text available: pdf(748.97 KB) Additional Information: full citation, abstract, references, index terms

Making high-bandwidth Internet access pervasively available to a large worldwide audience is a difficult challenge, especially in many developing regions. As we wait for the uncertain takeoff of technologies that promise to improve the situation, we propose to explore an approach that is potentially more easily realizable: the use of digital storage media transported by the postal system as a general digital communication mechanism. We shall call such a system a *Postmanet*. Compared to mor ...

Keywords: Distributed systems, peer-to-peer systems, postal system, storage systems, the digital divide

10 The space shuttle primary computer system

Alfred Spector, David Gifford
September 1984 **Communications of the ACM**, Volume 27 Issue 9

Full text available: pdf(5.34 MB) Additional Information: full citation, references, citings, index terms

Keywords: PASS, avionics system, space shuttle

11 A reconfigurable hardware approach to network simulation

Dimitrios Stiliadis, Anujan Varma
January 1997 **ACM Transactions on Modeling and Computer Simulation (TOMACS)**, Volume 7 Issue 1

Full text available: pdf(925.18 KB) Additional Information: full citation, references, citings, index terms, review

Keywords: ATM switch scheduling, field-programmable gate array, hardware simulation

12 InfiniteReality: a real-time graphics system
John S. Montrym, Daniel R. Baum, David L. Dignam, Christopher J. Migdal
August 1997 **Proceedings of the 24th annual conference on Computer graphics and interactive techniques**
Full text available: pdf(697.27 KB) Additional Information: full citation, references, citings, index terms

13 S-connect: from networks of workstations to supercomputer performance
Andreas G. Nowatzyk, Michael C. Browne, Edmund J. Kelly, Michael Parkin
May 1995 **ACM SIGARCH Computer Architecture News , Proceedings of the 22nd annual international symposium on Computer architecture**, Volume 23 Issue 2
Full text available: pdf(1.38 MB) Additional Information: full citation, abstract, references, citings, index terms

S-Connect is a new high speed, scalable interconnect system that has been developed to support networks of workstations to efficiently share computing resources. It uses off-the-shelf CMOS technology to directly drive fiber-optic systems at speeds greater than 1 Gbit/sec and can realize bisection bandwidths comparable to high-end MPP systems while being >10x more cost-effective. S-Connect systems do not rely on centralized switches, but rather are composed of adaptive, topology independen ...

14 Are crossbars really dead?: the case for optical multiprocessor interconnect systems
Andreas G. Nowatzyk, Paul R. Prucnal
May 1995 **ACM SIGARCH Computer Architecture News , Proceedings of the 22nd annual international symposium on Computer architecture**, Volume 23 Issue 2
Full text available: pdf(1.16 MB) Additional Information: full citation, abstract, references, citings, index terms

Crossbar switches are rarely considered for large, scalable multiprocessor interconnect systems because they require $O(n^2)$ switching elements, are difficult to control efficiently and are hard to implement once their size becomes too large to fit on one integrated circuit. However these problems are technology dependent and a recent innovation in fiber optic devices has led to a new implementation of crossbar switches that does not share these problems while retaining the full advanta ...

15 The Mercury Interconnect Architecture: a cost-effective infrastructure for high-performance servers
Wolf-Dietrich Weber, Stephen Gold, Pat Helland, Takeshi Shimizu, Thomas Wicki, Winfried Wilcke
May 1997 **ACM SIGARCH Computer Architecture News , Proceedings of the 24th annual international symposium on Computer architecture**, Volume 25 Issue 2
Full text available: pdf(1.53 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents HAL's Mercury Interconnect Architecture, an interconnect infrastructure designed to link commodity microprocessors, memory, and I/O components into high-performance multiprocessing servers. Both shared-memory and message-passing systems, as well as hybrid systems are supported by the interconnect. The key attributes of the Mercury Interconnect Architecture are: low latency, high bandwidth, a modular and flexible design, reliability/availability/serviceability (RAS) features, ...

16 System support for pervasive applications
Robert Grimm, Janet Davis, Eric Lemar, Adam Macbeth, Steven Swanson, Thomas Anderson, Brian Bershad, Gaetano Borriello, Steven Gribble, David Wetherall

November 2004 **ACM Transactions on Computer Systems (TOCS)**, Volume 22 Issue 4

Full text available: pdf(1.82 MB) Additional Information: full citation, abstract, references, index terms

Pervasive computing provides an attractive vision for the future of computing. Computational power will be available everywhere. Mobile and stationary devices will dynamically connect and coordinate to seamlessly help people in accomplishing their tasks. For this vision to become a reality, developers must build applications that constantly adapt to a highly dynamic computing environment. To make the developers' task feasible, we present a system architecture for pervasive computing, called & ...

Keywords: Asynchronous events, checkpointing, discovery, logic/operation pattern, migration, one.world, pervasive computing, structured I/O, tuples, ubiquitous computing

17 Design, implementation, and evaluation of a software-based real-time Ethernet protocol

Chitra Venkatramani, Tzi-cker Chiueh

October 1995 **ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication**, Volume 25 Issue 4

Full text available: pdf(1.18 MB) Additional Information: full citation, abstract, references, citings, index terms

Distributed multimedia applications require performance guarantees from the underlying network subsystem. Ethernet has been the dominant local area network architecture in the last decade, and we believe that it will remain popular because of its cost-effectiveness and the availability of higher-bandwidth Ethernets. We present the design, implementation and evaluation of a software-based timed-token protocol called RETHER that provides real-time performance guarantees to multimedia applications ...

18 The VMP network adapter board (NAB): high-performance network communication for multiprocessors

H. Kanakia, D. Cheriton

August 1988 **ACM SIGCOMM Computer Communication Review , Symposium proceedings on Communications architectures and protocols**, Volume 18 Issue 4

Full text available: pdf(1.63 MB) Additional Information: full citation, abstract, references, citings, index terms

High performance computer communication between multiprocessor nodes requires significant improvements over conventional host-to-network adapters. Current host-to-network adapter interfaces impose excessive processing, system bus and interrupt overhead on a multiprocessor host. Current network adapters are either limited in function, wasting key host resources such as the system bus and the processors, or else intelligent but too slow, because of complex transport protocols and because of a ...

19 Accelerating shared virtual memory via general-purpose network interface support

Angelos Bilas, Dongming Jiang, Jaswinder Pal Singh

February 2001 **ACM Transactions on Computer Systems (TOCS)**, Volume 19 Issue 1

Full text available: pdf(178.88 KB) Additional Information: full citation, abstract, references, index terms, review

Clusters of symmetric multiprocessors (SMPs) are important platforms for high-performance computing. With the success of hardware cache-coherent distributed shared memory (DSM), a lot of effort has also been made to support the coherent shared-address-space programming model in software on clusters. Much research has been done in fast communication on clusters and in protocols for supporting software shared memory across

them. However, the performance of software virtual memory (SVM) is sti ...

Keywords: applications, clusters, shared virtual memory, system area networks

20 A new method to make communication latency uniform: distributed routing balancing 
D. Franco, I. Garcés, E. Luque
May 1999 **Proceedings of the 13th international conference on Supercomputing**
Full text available: pdf(1.24 MB) Additional Information: full citation, references, citings, index terms

Keywords: adaptive routing, distributed routing balancing, hot spot avoidance, interconnection networks, random routing, traffic distribution, uniform latency

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next